Optimal Auctions Through Deep Learning

By Paul Dütting, Zhe Feng, Harikrishna Narasimhan, David C. Parkes, and Sai S. Ravindranath

Abstract

Designing an incentive compatible auction that maximizes expected revenue is an intricate task. The single-item case was resolved in a seminal piece of work by Myerson in 1981. Even after 30–40 years of intense research, the problem remains unsolved for settings with two or more items. We overview recent research results that show how tools from deep learning are shaping up to become a powerful tool for the automated design of near-optimal auctions. In this approach, an auction is modeled as a multi-layer neural network, with optimal auction design framed as a constrained learning problem that can be addressed with standard machine learning pipelines. Through this approach, it is possible to recover to a high degree of accuracy essentially all known analytically derived solutions for multi-item settings and obtain novel mechanisms for settings in which the optimal mechanism is unknown.

1. INTRODUCTION

Optimal auction design is one of the cornerstones of economic theory. It is of great practical importance, as auctions are used across industries and by the public sector to organize the sale of their products and services. Concrete examples are the US FCC Incentive Auction, the sponsored search auctions conducted by web search engines such as Google, and the auctions run on platforms such as eBay. In the standard *independent private valuations* model, each bidder has a valuation function over subsets of items, drawn independently from not necessarily identical distributions. The auctioneer knows the value distribution, but not the actual valuations (willingness to pay) of bidders. The bidders may act strategically and report untruthfully if this is to their benefit. One way to circumvent this is to require that it is in each agent's best interest to report its value truthfully. The goal then is to learn an incentive compatible auction that maximizes revenue.

In a seminal piece of work, Myerson resolved the optimal auction design problem when there is a single item for sale.¹⁷ Quite astonishingly, even after 30–40 years of intense research, the problem is not completely resolved even for a simple setting with two bidders and two items. Our focus is on designing auctions that satisfy *dominant-strategy incentive compatibility (DSIC)*, which is a robust and desirable notion of incentive alignment. Although there have been some elegant partial characterization results,^{6, 10, 15, 20} and an impressive sequence of algorithmic results, for example, Babaioff et al.¹ and Cai et al.,² these apply to the weaker notion of *Bayesian incentive compatibility (BIC)* except for the setting with one bidder, when DSIC and BIC coincide.

In Dütting et al.,⁷ we have introduced a new, deep-learning-based approach to address the problem of optimal, multi-item auction design. In particular, we use multilayer neural networks to encode auction mechanisms, with bidder valuations forming the input and allocation and payment decisions forming the output. The networks are trained using samples from the value distributions, so as to maximize expected revenue subject to constraints for incentive compatibility. Earlier work has suggested to use algorithms to automate the design of mechanisms,³ but where scalable, this earlier work had to restrict the search space to auction designs that are known to be incentive compatible.^{13, 23} The deep learning approach, in contrast, enables searching over broad classes of not necessarily truthful mechanisms. Another related line of work has leveraged machine learning to optimize different aspects of mechanisms,^{8, 18} but none of these offers the generality and flexibility of our approach.

Our framework provides two different approaches to handling DSIC constraints. In the first, we leverage results from economic theory that characterize DSIC mechanisms and model the network architecture appropriately. This approach, which we refer to as *RochetNet*, is applicable in single-bidder multi-item settings and provides exactly DSIC mechanisms.²² In the second, we lift the DSIC constraints into the objective via the *augmented Lagrangian* method, which has the effect of introducing a penalty term for DSIC violations. This approach, which we refer to as *RegretNet*, is also applicable in multibidder multi-item settings for which we do not have tractable characterizations of DSIC mechanisms but will generally only find mechanisms that are approximately incentive compatible.

In this Research Highlight, we describe the general approach and present a selection of experimental results in support of our general finding that these approaches are capable of recovering, to a high degree of accuracy, the optimal auctions from essentially all analytical results obtained over the past 30–40 years and that deep learning is also a powerful tool for confirming or refuting hypotheses concerning the form of optimal auctions and can be used to find new designs. In the full version of the paper, we also prove *generalization bounds* that provide confidence intervals on the expected revenue and expected violation of DSIC based on empirical properties obtained

An extended abstract was published in *Proceedings of the 36th International Conference on Machine Learning, 2019*. A full version of this paper is available at <https://arxiv.org/abs/1706.03459>. All code is available through the GitHub repository at <https://github.com/saisrivatsan/deep-opt-auctions>.

during training, the complexity of the neural network used to encode the allocation and payment rules, and the number of samples used to train the network. Others have provided generalization bounds for training revenue-maximizing auctions in simpler settings; see, for example, Morgenstern and Roughgarden.¹⁶

Follow-up work has extended our approach to handle *budget constraints*,⁹ as well as to a problem in social choice, the so-called *facility location problem*,¹² studied specialized architectures for single-bidder settings,²⁴ introduced networks that encode symmetry,²¹ and provided methods to certify the strategy-proofness of learned mechanisms.⁴

2. OPTIMAL AUCTION DESIGN

We start by stating the optimal auction design problem and providing a few illustrative examples.

In the general version of the problem, we are given n bidders $N = \{1, \dots, n\}$ and m items $M = \{1, \dots, m\}$. Each bidder i has a valuation function $v_i: 2^M \rightarrow \mathbb{R}_{\geq 0}$, where $v_i(S)$ denotes how much the bidder values the subset of items $S \subseteq M$. In the simplest case, a bidder may have *additive* valuations. In this case, she has a value $v_i(\{j\})$ for each individual item $j \in M$, and her value for a subset of items $S \subseteq M$ is $v_i(S) = \sum_{j \in S} v_i(\{j\})$. If a bidder's value for a subset of items $S \subseteq M$ is $v_i(S) = \max_{j \in S} v_i(\{j\})$, we say this bidder has a *unit-demand* valuation. We also consider bidders with specific combinatorial valuations but defer the details to our full version.

Bidder i 's valuation function is drawn independently from a distribution F_i over possible valuation functions V_i . We write $v = (v_1, \dots, v_n)$ for a profile of valuations and denote $V = \prod_{i=1}^n V_i$. The auctioneer knows the distributions $F = (F_1, \dots, F_n)$ but does not know the bidders' realized valuation v . The bidders report their valuations (perhaps untruthfully), and an auction decides on an allocation of items to the bidders and charges a payment to them. We denote an auction (g, p) as a pair of allocation rules $g_i: V \rightarrow 2^M$ and payment rules $p_i: V \rightarrow \mathbb{R}_{\geq 0}$ (these rules can be randomized). Given bids $b = (b_1, \dots, b_n) \in V$, the auction computes an allocation $g(b)$ and payments $p(b)$.

A bidder with valuation v_i receives a utility $u_i(v_i; b) = v_i(g_i(b)) - p_i(b)$ for a report of bid profile b . Let v_{-i} denote the valuation profile $v = (v_1, \dots, v_n)$ without element v_i , similarly for b_{-i} , and let $V_{-i} = \prod_{j \neq i} V_j$ denote the possible valuation profiles of bidders other than bidder i . An auction is *dominant strategy incentive compatible* (DSIC) if each bidder's utility is maximized by reporting truthfully no matter what the other bidders report. In other words, $u_i(v_i; (v_{-i}, b_{-i})) \geq u_i(v_i; (v_{-i}, b_{-i}'))$ for every bidder i , every valuation $v_i \in V_i$, every bid $b_{-i} \in V_{-i}$, and all bids $b_{-i}' \in V_{-i}$ from others. An auction is *ex post individually rational* (IR) if each bidder receives a nonzero utility, that is, $u_i(v_i; (v_{-i}, b_{-i})) \geq 0 \forall i \in N, v_i \in V_i$, and $b_{-i} \in V_{-i}$.

In a DSIC auction, it is in the best interest of each bidder to report truthfully, and so the revenue on valuation profile v is $\sum_i p_i(v)$. Optimal auction design seeks to identify a DSIC auction that maximizes expected revenue.

EXAMPLE 1 (VICKREY AUCTION²⁶). *A classic result in auction theory concerns the sale of a single item to n bidders. It states that the following auction—the so-called Vickrey or second-price*

auction—is DSIC and maximizes social welfare: Collect a bid b_i from each bidder, assign the item to the bidder with the highest bid (breaking ties in an arbitrary but fixed manner), and make the bidder pay the second-highest bid.

EXAMPLE 2 (MYERSON AUCTION¹⁷). *A simple example shows that the Vickrey auction does not maximize revenue: Suppose there are two bidders with $v_i \in U[0, 1]$, then its expected revenue is $1/3$. Higher revenue can be achieved with a second-price auction with reserve r : As before, collect bids b_i , allocate to the highest bid but only if this bid is at least r , and make the winning bidder (if any) pay the maximum of the runner-up bid and r . It is straightforward to verify that this auction is DSIC and that choosing $r = 1/2$ leads to an expected revenue of $5/12 > 1/3$.*

In the simple example with a single item and uniform valuations, a second-price auction with reserve $1/2$ is in fact the optimal auction. This auction illustrates a special case of Myerson's theory for the design of revenue-optimal, single-item auctions.¹⁷ Comparable results are not available for selling multiple items, even when we are trying to sell them to a single bidder!

3. THE LEARNING PROBLEM

At the core of our approach is the following reinterpretation of the optimal auction design problem as a learning problem, where in the place of a loss function that measures error against a target label, we adopt the negated, expected revenue on valuations drawn from F .

More concretely, the problem we seek to solve is the following: We are given a parametric class of auctions, $(g^w, p^w) \in \mathcal{M}$, for parameters $w \in \mathcal{R}^d$ for some $d \in \mathcal{N}$, and a sample of bidder valuation profiles $\mathcal{S} = \{v^{(1)}, \dots, v^{(L)}\}$ drawn i.i.d. from F . Our goal is to find an auction that minimizes the negated, expected revenue $-\mathbb{E}[\sum_{i \in N} p_i^w(v)]$ among all auctions in $\mathcal{M}$ that satisfy incentive compatibility.

We consider two distinct approaches for achieving DSIC. In the first approach, we make use of characterization results. When it is possible to encode them within a neural network architecture, these characterizations from economic theory usefully constrain the search space and provide exact DSIC. At the same time, the particular characterization that we use is limited in that it applies only to single-bidder settings. The second approach that we take is more general, applying to multi-bidder settings, and does not rely on the availability of suitable characterization results. On the other hand, this approach entails search through a larger parametric space and only achieves approximate DSIC.

We describe the first approach in Section 4 and return to the second approach in Section 5.

4. THE ROCHETNET FRAMEWORK

We have developed two different frameworks that achieve exact DSIC by applying appropriate structure to the neural network architecture. One framework, referred to as *MyersonNet*, is inspired by Myerson's lemma¹⁷ and can be used for the study of multi-bidder, single-item auctions (see the full version of this paper). A second framework,

referred to as *RochetNet*, is inspired by Rochet’s characterization theorem for DSIC auctions in single-bidder settings.²² We give the construction of RochetNet for additive preferences, but this can be easily extended to unit-demand valuations.

4.1. The RochetNet architecture

For this single-bidder, multi-item setting, let $v \in \mathbb{R}_{\geq 0}^m$ denote the bidder’s additive valuation, so that v_j is its value for item j . Let $b \in \mathbb{R}_{\geq 0}^m$ denote the bid, which need not be truthful. The allocation rule $g^w : \mathbb{R}_{\geq 0}^m \rightarrow [0, 1]^m$, for parameters w , defines for each item $j \in [J]$ the probability $g_j^w(b) \in [0, 1]$ with which the item is allocated to the bidder. The payment rule $p^w(b) : \mathbb{R}_{\geq 0}^m \rightarrow \mathbb{R}$ defines the payment $p^w(b)$ made by the bidder.

The mechanism (g^w, p^w) induces a *utility function* $u^w : \mathbb{R}_{\geq 0}^m \rightarrow \mathbb{R}$. For truthful bids, v , the utility function induced by the mechanism is

$$u^w(v) = g^w(v) \cdot v - p^w(v). \quad (1)$$

The RochetNet architecture represents the rules of a mechanism through a *menu*. The menu encodes a set of K choices, where each choice consists of a randomized allocation together with a price. The network selects the choice for the bidder that maximizes the bidder’s reported utility given its bid, or chooses the *null outcome* (no allocation, no payment) when this is preferred. This yields the following utility function:

$$u^w(v) = \max \left\{ \max_{k \in [K]} \{ \alpha_k \cdot v + \beta_k \}, 0 \right\}, \quad (2)$$

with parameters $w = (\alpha, \beta)$, where $\alpha \in [0, 1]^{mK}$ and $\beta \in \mathbb{R}^K$. For choice $k \in [K]$, parameters $\alpha_k \in [0, 1]^m$ specify the randomized allocation and parameter $\beta_k \in \mathbb{R}$ specifies the negated price (β_k s will be negative, and the smaller the value of β_k , the larger the payment).

For input b , let $k^*(b) \in \arg\max_{k \in [K] \cup \{0\}} \{ \alpha_k \cdot b + \beta_k \}$ denote the best choice for the bidder, where choice 0 corresponds to $\alpha_0 = 0$ and $\beta_0 = 0$ and the null outcome. This best choice defines the allocation and payment rule—for bid b , the allocation is $g^w(b) = \alpha_{k^*(b)}$ and the payment is $p^w(b) = -\beta_{k^*(b)}$.

RochetNet represents this induced utility function as a single layer neural network as illustrated in Figure 1(a). The input layer takes a bid $b \in \mathbb{R}_{\geq 0}^m$ and the output of the network

Figure 1. RochetNet: (a) Neural network representation of a menu, shown here with K choices as well as the null outcome (0); here, $h_k(b) = \alpha_k \cdot b + \beta_k$ for $b \in \mathbb{R}^m$, $\alpha_k \in [0, 1]^m$, and $\beta_k \in \mathbb{R}$. (b) An induced utility function represented by RochetNet for the case of a single item ($m = 1$) and a network with a menu with four choices ($K = 4$).

is the induced utility. Figure 1(b) shows an example of an induced utility function for a single item ($m = 1$) and a network with a menu consisting of four choices ($K = 4$).

The network architecture ensures that the utility function is monotonically non decreasing, convex, and 1-Lipschitz, conforming to Rochet’s characterization.²² It also easily provides the following theoretical property.

THEOREM 4.1. *For any parameterization w , the mechanism (g^w, p^w) corresponding to RochetNet is DSIC and IR.*

PROOF. For DSIC, note that (1) the available choices are fixed, and independent of the report; and (2) for a truthful report, the “max” structure of RochetNet ensures that the bidder receives the choice that maximizes its true expected utility, and thus, the bidder can do no better than this. For IR, note that the expected utility for a true report is at least zero because of the availability of the null outcome. $\square$

4.2. Training

During training, we seek to minimize the negated, expected revenue. Let F denote the distribution on valuation v . To ensure that the objective is a continuous function of α and β (so that parameters can be optimized through gradient descent), the best choice $k^*(v)$ at input v is approximated during training via a *softmax* operation in place of the argmax. With this, we seek to minimize the following loss function, which corresponds to the approximate, negated revenue:

$$\mathcal{L}(\alpha, \beta) = \mathbb{E}_{v \sim F} \left[\sum_{k \in [K] \cup \{0\}} \beta_k x_k(v) \right], \quad (3)$$

where

$$x_{k(v)} = \text{softmax}_k(0, c(\alpha_1 \cdot v + \beta_1), \dots, c(\alpha_K \cdot v + \beta_K)) \quad (4)$$

and $c > 0$ is a constant that controls the quality of the approximation. The softmax function is $\text{softmax}_k(cz_0, cz_1, \dots, cz_K) = e^{cz_k} / \sum_{k'} e^{cz_{k'}}$ and takes as input $K + 1$ real numbers and returns a probability distribution with each entry proportional to the exponential of the corresponding input. Once trained, RochetNet is used at test time with a hard max in place of the softmax to ensure exact DSIC and IR.

We train RochetNet using samples drawn from the bidder’s value distribution. Given a sample $\mathcal{S} = \{v^{(1)}, \dots, v^{(L)}\}$, we minimize the empirical loss, which is

$$\min_{\alpha, \beta} \frac{1}{L} \sum_{\ell \in [L]} \left(\sum_{k \in [K] \cup \{0\}} \beta_k x_k(v^{(\ell)}) \right). \quad (5)$$

We use projected stochastic gradient descent (SGD) to minimize (5). We estimate gradients for the loss using mini-batches of size 2^{15} valuation samples in every iteration. In the projection step, we project each parameter α_{jk} (for item j , choice k) onto $[0, 1]$ to provide a well-defined probability.

5. THE REGRETNET FRAMEWORK

We next describe our second approach to handling DSIC constraints and the corresponding framework, which we refer to as *RegretNet*. Unlike the first approach, this second

approach does not rely on characterizations of DSIC mechanisms. Instead, we replace the DSIC constraints with a differentiable approximation and lift the DSIC constraints into the objective by augmenting the objective with a term that accounts for the extent to which the DSIC constraints are violated. Here, we provide an overview of the special case in which bidders have additive values for items, but the framework also handles more general settings.

5.1. Expected ex post regret

We can measure the extent to which an auction violates incentive compatibility through a particular variation on *ex post* regret introduced in Dütting et al.⁸ Fixing the bids of others, the *ex post* regret for a bidder is the maximum increase in her utility, considering all possible nontruthful bids.

For mechanisms (g^w, p^w) , we will be interested in the *expected ex post regret* for bidder i :

$$rgt_i(w) = \mathbb{E} \left[\max_{v'_i \in V_i} u_i^w(v_i; (v'_i, v_{-i})) - u_i^w(v_i; (v_i, v_{-i})) \right],$$

where the expectation is over $v \sim F$ and $u_i^w(v_i; b) = v_i(g_i^w(b)) - p_i^w(b)$ for model parameters w . We assume that F has full support on the space of valuation profiles V , and recognizing that the regret is nonnegative, an auction satisfies DSIC if and only if $rgt_i(w) = 0, \forall i \in N$, except for measure zero events.

Given this, we reformulate the learning problem as minimizing expected negated revenue subject to the expected *ex post* regret being zero for each bidder:

$$\begin{aligned} \min_{w \in \mathbb{R}^d} \quad & \mathbb{E}_{v \sim F} \left[-\sum_{i \in N} p_i^w(v) \right] \\ \text{s.t.} \quad & rgt_i(w) = 0, \forall i \in N. \end{aligned}$$

Given a sample $\mathcal{S}$ of L valuation profiles from F , we estimate the empirical *ex post* regret for bidder i as:

$$\widehat{rgt}_i(w) = \frac{1}{L} \sum_{\ell=1}^L \left[\max_{v'_i \in V_i} u_i^w(v_i^{(\ell)}; (v'_i, v_{-i}^{(\ell)})) - u_i^w(v_i^{(\ell)}; v_{-i}^{(\ell)}) \right], \quad (6)$$

and seek to minimize the empirical loss (negated revenue) subject to the empirical regret being zero for all bidders:

$$\begin{aligned} \min_{w \in \mathbb{R}^d} \quad & -\frac{1}{L} \sum_{\ell=1}^L \sum_{i=1}^n p_i^w(v^{(\ell)}) \\ \text{s.t.} \quad & \widehat{rgt}_i(w) = 0, \forall i \in N. \end{aligned} \quad (7)$$

We additionally require the designed auction to satisfy IR, which can be ensured by restricting the search space to a class of parameterized auctions that charge no bidder more than her valuation for an allocation.

5.2. The RegretNet architecture

In this case, the goal is to train neural networks that explicitly encode the allocation and payment rule of the mechanism. The architectures generally consist of two logically distinct components: the allocation and payment networks. These components are trained together and the outputs of these networks are used to compute the regret and revenue of the auction.

An overview of the RegretNet architecture for additive valuations is given in Figure 2.

The allocation network encodes a randomized allocation rule $g^w : \mathbb{R}^{nm} \rightarrow [0, 1]^{nm}$ and the payment network encodes a payment rule $p^w : \mathbb{R}^{nm} \rightarrow \mathbb{R}_{\geq 0}^m$, both of which are modeled as feedforward, fully-connected networks with a tanh activation function in each of the hidden nodes. The input layer of the networks consists of bids $b_{ij} \geq 0$ representing the valuation of bidder i for item j .

The allocation network outputs a vector of allocation probabilities $z_{1j} = g_{1j}(b), \dots, z_{nj} = g_{nj}(b)$, for each item $j \in [m]$. To ensure feasibility, that is, the probability of an item being allocated is at most one, the allocations are computed using a softmax activation function, so that for all items j , we have $\sum_{i=1}^n z_{ij} \leq 1$. To accommodate the possibility of an item not being assigned, we include a dummy node in the softmax computation to hold the residual allocation probability. The payment network outputs a payment for each bidder that denotes the amount the bidder should pay in expectation for a particular bid profile.

To ensure that the auction satisfies IR, that is, does not charge a bidder more than her expected value for the

Figure 2. RegretNet: The allocation and payment networks for a setting with n additive bidders and m items. The inputs are bids from each bidder for each item. The revenue rev and expected *ex post* rgt_i are defined as a function of the parameters of the allocation and payment networks $w = (w_g, w_p)$.

allocation, the network first computes a normalized payment $\tilde{p}_i \in [0, 1]$ for each bidder i using a sigmoidal unit, and then outputs a payment $p_i = \tilde{p}_i (\sum_{j=1}^m z_{ij} b_{ij})$, where the z_{ij} 's are the outputs from the allocation network.

5.3. Training

For RegretNet, we have used the *augmented Lagrangian method* to solve the constrained training problem in (7) over the space of neural network parameters w .

Algorithm 1 RegretNet Training

```

1: Input: Minibatches  $\mathcal{S}_1, \dots, \mathcal{S}_T$  of size  $B$ 
2: Parameters:  $\forall t, \rho_t > 0, \gamma > 0, \eta > 0, \Gamma \in \mathbb{N}, K \in \mathbb{N}$ 
3: Initialize:  $w^0 \in \mathbb{R}^d, \lambda^0 \in \mathbb{R}^n$ 
4: for  $t = 0$  to  $T$  do
5:   Receive minibatch  $\mathcal{S}_t = \{v^{(1)}, \dots, v^{(B)}\}$ 
6:   Initialize misreports  $v_i^{(\ell)} \in V_i, \forall \ell \in [B], i \in N$ 
7:   for  $r = 0$  to  $\Gamma$  do
8:      $\forall \ell \in [B], i \in N$ :
9:      $v_i^{(\ell)} \leftarrow v_i^{(\ell)} + \gamma \nabla_{v_i} u_i^w(v_i^{(\ell)}; (v_i^{(\ell)}, v_{-i}^{(\ell)}))$ 
10:  end for
11:  Compute regret gradient:  $\forall \ell \in [B], i \in N$ :
12:   $g_{\ell,i}^t = \nabla_w [u_i^w(v_i^{(\ell)}; (v_i^{(\ell)}, v_{-i}^{(\ell)})) - u_i^w(v_i^{(\ell)}; v^{(\ell)})]_{w=w^t}$ 
13:  Compute Lagrangian gradient (8) on  $\mathcal{S}_t$  and update:
14:   $w^{t+1} \leftarrow w^t - \eta \nabla_w \mathcal{C}_\rho(w^t, \lambda^t)$ 
15:  Update Lagrange multipliers once in  $Q$  iterations:
16:  if  $t$  is a multiple of  $Q$ 
17:    Compute  $\widehat{rgt}$  on  $\mathcal{S}_t$ 
18:     $\lambda_i^{t+1} \leftarrow \lambda_i^t + \rho_i \widehat{rgt}_i(w^{t+1}), \forall i \in N$ 
19:  else
20:     $\lambda^{t+1} \leftarrow \lambda^t$ 
21:  end for
22: end for

```

We first define the Lagrangian function for the optimization problem, augmented with a quadratic penalty term for violating the constraints:

$$\mathcal{C}_\rho(w; \lambda) = -\frac{1}{L} \sum_{\ell=1}^L \sum_{i \in N} p_i^w(v^{(\ell)}) + \sum_{i \in N} \lambda_i \widehat{rgt}_i(w) + \frac{\rho}{2} \sum_{i \in N} (\widehat{rgt}_i(w))^2$$

where $\lambda \in \mathbb{R}^n$ is a vector of Lagrange multipliers and $\rho > 0$ is a fixed parameter that controls the weight on the quadratic penalty. The solver alternates between the following updates on the model parameters and the Lagrange multipliers: (a) $w^{new} \in \argmin_w \mathcal{C}_\rho(w^{old}; \lambda^{old})$ and (b) $\lambda_i^{new} = \lambda_i^{old} + \rho \widehat{rgt}_i(w^{new}), \forall i \in N$.

The solver is described in Algorithm 1. We divide the training sample $\mathcal{S}$ into minibatches of size B , estimate gradients on the minibatches, and perform several passes over the training samples. The update (a) on model parameters involves an unconstrained optimization of $\mathcal{C}_\rho$ over w and is performed using a gradient-based optimizer. The gradient $\mathcal{C}_\rho$ of w.r.t. w for fixed λ^t is given by:

$$\begin{aligned} \nabla_w \mathcal{C}_\rho(w; \lambda^t) = & -\frac{1}{B} \sum_{\ell=1}^B \sum_{i \in N} \nabla_w p_i^w(v^{(\ell)}) + \sum_{i \in N} \sum_{\ell=1}^B \lambda_i^t g_{\ell,i} \\ & + \rho \sum_{i \in N} \sum_{\ell=1}^B \widehat{rgt}_i(w) g_{\ell,i}, \end{aligned} \quad (8)$$

where

$$g_{\ell,i} = \nabla_w \left[\max_{v_i' \in V_i} u_i^w(v_i^{(\ell)}; (v_i', v_{-i}^{(\ell)})) - u_i^w(v_i^{(\ell)}; v^{(\ell)}) \right].$$

The terms $\widehat{rgt}_i$ and $g_{\ell,i}$ in turn involve a “max” over misreports for each bidder i and valuation profile ℓ . We solve this inner maximization over misreports using another gradient-based optimizer (lines 6–10).

As the optimization problem is nonconvex, the solver is not guaranteed to reach a globally optimal solution. However, this method proves very effective in our experiments, and we find that the learned auctions incur very low regret and closely match the structure of optimal auctions in settings where this is known.

6. EXPERIMENTS

We present and discuss a selection of experiments out of a broad range of experiments that we have conducted and that we describe in more detail in Dütting et al.⁷ and the full version. The experiments demonstrate that our approach can recover near-optimal auctions for essentially all settings for which the optimal design is analytically known, that it is an effective tool for confirming or refuting hypotheses about optimal designs, and that it can find new auctions for settings where there is no known analytical solution.

6.1. Setup

We implemented our framework using the TensorFlow deep learning library.

For *RochetNet*, we initialized parameters α and β in Eq. (2) using a random uniform initializer over the interval $[0, 1]$ and a zero initializer, respectively. For *RegretNet*, we used the tanh activation function at the hidden nodes, and Glorot uniform initialization.¹¹ We performed cross-validation to decide on the number of hidden layers and the number of nodes in each hidden layer. We include exemplary numbers that illustrate the trade-offs in Section 6.6.

We trained *RochetNet* on 2^{15} valuation profiles and sampled every iteration in an online manner. We used the Adam optimizer with a learning rate of 0.1 for 20,000 iterations for making the updates. The parameter κ in Eq. (4) was set to 1000. Unless specified otherwise, we used a max network over 1000 linear functions to model the induced utility functions and report our results on a sample of 10,000 profiles.

For *RegretNet*, we used a sample of 640,000 valuation profiles for training and a sample of 10,000 profiles for testing. The augmented Lagrangian solver was run for a maximum of 80 epochs (full passes over the training set) with a minibatch size of 128. The value of ρ in the augmented Lagrangian was set to 1.0 and incremented every two epochs. An update on w^t was performed for every minibatch using the Adam optimizer with learning rate 0.001. For each update on w^t , we ran $\Gamma = 25$ misreport update steps with learning rate 0.1. At the end of 25 updates, the optimized misreports for the current minibatch were cached and used to initialize the misreports for the same minibatch in the next epoch. An update on λ^t was performed once every 100 minibatches (i.e., $Q = 100$).

We ran all our experiments on a compute cluster with NVIDIA Graphics Processing Unit (GPU) cores.

6.2. Evaluation

In addition to the revenue of the learned auction on a test set, we also evaluate the regret achieved by RegretNet, averaged across all bidders and test valuation profiles, that is, $rgt = \frac{1}{n} \sum_{i=1}^n \widehat{rgt}_i(g^w, p^w)$. Each $\widehat{rgt}_i$ has an inner “max” of the utility function over bidder valuations $v'_i \in V_i$ (see (6)). We evaluate these terms by running gradient ascent on v'_i with a step-size of 0.1 for 2000 iterations (we test 1000 different random initial v'_i and report the one that achieves the largest regret). For some of the experiments, we also report the total time it took to train the network. This time is incurred during offline training, whereas the allocation and payments can be computed in a few milliseconds once the network is trained.

6.3. The Manelli-Vincent auction

As a representative example the optimal designs from economic theory that we can almost exactly recover with our approach, we discuss the Manelli-Vincent auction.¹⁵

- A. Single bidder with additive valuations over two items, where the item values are independent draws from $U[0, 1]$.

The optimal auction for this setting is given by Manelli and Vincent.¹⁵ We used two hidden layers with 100 hidden nodes in RegretNet for this setting. A visualization of the optimal allocation rule and those learned by RocketNet and RegretNet is given in Figure 3. Figure 4(a) gives the optimal revenue, the revenue and regret obtained by RegretNet, and the revenue obtained by RocketNet. Figure 4(b) shows how these terms evolve over time during training in RegretNet.

Both approaches essentially recover the optimal design, not only in terms of revenue but also in terms of the allocation rule and transfers. The auction learned by RocketNet is exactly DSIC and matches the optimal revenue precisely, with sharp decision boundaries in the allocation and payment rule. The decision boundaries for RegretNet are smoother, but still remarkably accurate. The revenue achieved by RegretNet matches the optimal revenue up to a <1% error term and the regret it incurs is <0.001. The plots of the test revenue and regret show that the augmented Lagrangian method is effective in driving the test revenue and the test regret toward optimal levels.

The additional domain knowledge incorporated into the RocketNet architecture leads to exactly DSIC mechanisms that match the optimal design more accurately and speeds up computation (the training took about 10 minutes compared to 11 hours for RegretNet). On the other hand, we find it surprising how well RegretNet performs given that it starts with no domain knowledge at all.

6.4. The Straight-Jacket auction

Extending the analytical result of Manelli and Vincent¹⁵ to a single bidder and an arbitrary number of items (even with additive preferences, all uniform on $[0, 1]$) has proven elusive. It is not even clear whether the optimal mechanism is deterministic or requires randomization.

Figure 3. Side-by-side comparison of allocation rules learned by RocketNet (panels (a)) and RegretNet (panels (b)) for Setting A. The panels describe the probability that the bidder is allocated item 1 (left) and item 2 (right) for different valuation inputs. The optimal auctions are described by the regions separated by the dashed black lines, with the numbers in black being the optimal probability of allocation in the region.

Figure 4. (a) Test revenue and regret for RegretNet and revenue for RocketNet for Setting A. (b) Plot of test revenue and regret as a function of training epochs for Setting A with RegretNet.

Distribution	Opt	RegretNet		RocketNet
		rev	rgt	rev
Setting A	0.550	0.554	<0.001	0.550

Giannakopoulos and Koutsoupias¹⁰ proposed a Straight-Jacket Auction (SJA) and gave a recursive algorithm for finding the subdivision and the prices, and used LP duality to prove that the SJA is optimal for items. These authors also conjecture that the SJA remains optimal for $m \leq 6$ general m but were unable to prove it.

Figure 5 gives the revenue of the SJA and that found by RocketNet for $m \leq 10$ items. We used a test sample of 2^{30} valuation profiles (instead of 10,000) to compute these numbers for higher precision. It shows that RocketNet finds the optimal revenue for $m \leq 6$ items and that it finds DSIC auctions whose revenue matches that of the SJA for $m = 7, 8, 9$, and 10 items. Closer inspection reveals that the allocation and payment rules learned by RocketNet essentially match those predicted by Giannakopoulos and Koutsoupias¹⁰ for all $m \leq 10$. We take this as strong additional evidence that the conjecture of Giannakopoulos and Koutsoupias¹⁰ is correct.

6.5. Discovering new optimal designs

RochetNet can also be used to aid the discovery of new, provably optimal designs. For this, we consider a single bidder with additive but correlated valuations for two items as follows:

- B. One additive bidder and two items, where the bidder's valuation is drawn uniformly from the triangle $T = \{(v_1, v_2) \mid \frac{v_1}{c} + v_2 \leq 2, v_1 \geq 0, v_2 \geq 1\}$ where $c > 0$ is a free parameter.

There is no analytical result for the optimal auction design for this setting. We ran RochetNet for different values of c (e.g., 0.5, 1, 3, 5) to discover the optimal auction. Based on this, we conjectured that the optimal mechanism contains two menu items for $c \leq 1$, namely $\{(0, 0), 0\}$ and $\{(1, 1), \frac{2+\sqrt{1+3c}}{3}\}$, and three menu items for $c > 1$, namely $\{(0, 0), 0\}$, $\{(1/c, 1), 4/3\}$, and $\{(1, 1), 1 + c/3\}$, giving the optimal allocation and payment in each region. In particular, as c transitions from values less than or equal to 1 to values larger than 1, the optimal mechanism transitions from being deterministic to being randomized. We have used duality theory⁵ to prove the optimality of this design, as stated in Theorem 6.1.

THEOREM 6.1. *For any $c > 0$, suppose the bidder's valuation is uniformly distributed over set $T = \{(v_1, v_2) \mid \frac{v_1}{c} + v_2 \leq 2, v_1 \geq 0, v_2 \geq 1\}$. Then, the optimal auction contains two menu items $\{(0, 0), 0\}$ and $\{(1, 1), \frac{2+\sqrt{1+3c}}{3}\}$ when $c \leq 1$, and three menu items $\{(0, 0), 0\}$, $\{(1/c, 1), 4/3\}$, and $\{(1, 1), 1 + c/3\}$ otherwise.*

6.6. Scaling up

We have also considered settings with up to five bidders and up to ten items. This is several orders of magnitude more complex than settings that can be addressed through other computational approaches to DSIC auction design. It is also a natural playground for RegretNet as no tractable characterizations of DSIC mechanisms are known for these settings.

The following two settings generalize the basic setting considered in Manelli and Vincent¹⁵ and Giannakopoulos and Koutsoupias¹⁰ to more than one bidder:

Figure 5. Revenue of the Straight-Jacket Auction (SJA) computed via the recursive formula in Giannakopoulos and Koutsoupias¹⁰ and that of the auction learned by RochetNet, for various numbers of items m . The SJA is known to be optimal for up to six items and conjectured to be optimal for any number of items.

Items	SJA (rev)	RochetNet (rev)
2	0.549187	0.549175
3	0.875466	0.875464
4	1.219507	1.219505
5	1.576457	1.576455
6	1.943239	1.943216
7	2.318032	2.318032
8	2.699307	2.699305
9	3.086125	3.086125
10	3.477781	3.477722

- C. Three additive bidders and ten items, where bidders draw their value for each item independently from the uniform distribution $U[0,1]$.
- D. Five additive bidders and ten items, where bidders draw their value for each item independently from the uniform distribution $U[0,1]$.

The optimal auction for these settings is not known. However, running a separate Myerson auction for each item is optimal in the limit of the number of bidders.¹⁹ For a regime with a small number of bidders, this provides a strong benchmark. We also compare to selling the grand bundle via a Myerson auction.

For Setting C, we show in Figure 6(a) the revenue and regret of the learned auction on a validation sample of 10,000 profiles, obtained with different architectures. Here, (R, K) denotes an architecture with R hidden layers and K nodes per layer. The (5, 100) architecture has the lowest regret among all the 100-node networks for both Setting C and Setting D. Figure 6(b) shows that the learned auctions yield higher revenue compared to the baselines, and do so with tiny regret.

7. CONCLUSION

The results from this research demonstrate that the methods of deep learning can be used to find close approximations to optimal designs from auction theory where they are known, to aid with the discovery of new optimal designs, and to scale-up computational approaches to optimal, DSIC auction design. Although our approach can be applied to settings that are orders of magnitude more complex than those that can be reached through other approaches to optimal DSIC design, a natural next step would be to scale this approach further to industry scale (e.g., through standardized benchmarking suites and innovations in network architecture). We also see promise for this framework in advancing economic theory, for example in supporting or refuting conjectures and as an assistant in guiding new economic discovery.

More generally, we believe that our work (together with a handful of contemporary works such as Hartford et al.,¹⁴ Thompson et al.²⁵) has opened the door to ML-assisted economic theory and practice, and we are looking forward to the advances that this agenda will bring along. □

Figure 6. (a) Revenue and regret of RegretNet on the validation set for auctions learned for Setting C using different architectures, where (R, K) denotes R hidden layers and K nodes per layer. (b) Test revenue and regret for Settings C and D, for the (5, 100) architecture.

Setting	RegretNet rev	RegretNet rgt	Item-wise Myerson	Bundled Myerson
C: 3×10	5.541	< 0.002	5.310	5.009
D: 5×10	6.778	< 0.005	6.716	5.453

(b)

References

- Babaioff, M., Immorlica, N., Lucier, B., Weinberg, S.M. A simple and approximately optimal mechanism for an additive buyer. In *Proceedings of the 55th IEEE Symposium on Foundations of Computer Science*, 2014, 21–30.
- Cai, Y., Daskalakis, C., Weinberg, S.M. An algorithmic characterization of multi-dimensional mechanisms. In *Proceedings of the 44th ACM Symposium on Theory of Computing*, 2012, 459–478.
- Conitzer, V., Sandholm, T. Complexity of mechanism design. In *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence*, 2002, 103–110.
- Curry, M.J., Chiang, P.-Y., Goldstein, T., Dickerson, J.P. Certifying strategyproof auction network. In *Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS 2020)*.
- Daskalakis, C., Deckelbaum, A., Tzamos, C. Mechanism design via optimal transport. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, 2013, 269–286.
- Daskalakis, C., Deckelbaum, A., Tzamos, C. Strong duality for a multiple-good monopolist. *Econometrica*, 85 (2017), 735–767.
- Dütting, P., Feng, Z., Narasimhan, H., Parkes, D.C., Ravindranath, S.S. Optimal auctions through deep learning. In *Proceedings of the 36th International Conference on Machine Learning*, 2019, 1706–1715.
- Dütting, P., Fischer, F., Jirapinyo, P., Lai, J., Lubin, B., Parkes, D.C. Payment rules through discriminant-based classifiers. *ACM Trans. Econ. Comput.* 1, 3 (2014), 5.
- Feng, Z., Narasimhan, H., Parkes, D.C. Deep learning for revenue-optimal auctions with budgets. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems*, 2018, 354–362.
- Giannakopoulos, Y., Koutsopoulos, E. Duality and optimality of auctions for uniform distributions. In *SIAM J. Comput.*, 47 (2018), 121–165.
- Glorot, X., Bengio, Y. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics*, 2010.
- Golowich, N., Narasimhan, H., Parkes, D.C. Deep learning for multi-facility location mechanism design. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, 2018, 261–267.
- Guo, M., Conitzer, V. Computationally feasible automated mechanism design: General approach and case studies. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, 2010.
- Hartford, J.S., Wright, J.R., Leyton-Brown, K. Deep learning for predicting human strategic behavior. In *Proceedings of the 29th Conference on Neural Information Processing Systems*, 2016, 2424–2432.
- Manelli, A., Vincent, D. Bundling as an optimal selling mechanism for a multiple-good monopolist. *J. Econ. Theory* 1, 127 (2006), 1–35.
- Morgenstern, J., Roughgarden, T. On the pseudo-dimension of nearly optimal auctions. In *Proceedings of the 28th Conference on Neural Information Processing Systems*, 2015, 136–144.
- Myerson, R. Optimal auction design. *Math. Operat. Res.*, 6 (1981), 58–73.
- Narasimhan, H., Agarwal, S., Parkes, D.C. Automated mechanism design without money via machine learning. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, 2016, 433–439.
- Palfrey, T. Bundling decisions by a multiproduct monopolist with incomplete information. *Econometrica*, 2, 51 (1983), 463–483.
- Pavlov, G. Optimal mechanism for selling two goods. *B.E. J. Theor. Econ.*, 11 (2011), 1–35.
- Rahme, J., Jelassi, S., Bruna, J., Weinberg, S.M. A permutation-equivariant neural network architecture for auction design. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence* (2021).
- Rochet, J.-C. A necessary and sufficient condition for rationalizability in a quasilinear context. *J. Math. Econ.*, 16 (1987), 191–200.
- Sandholm, T., Likhodedov, A. Automated design of revenue-maximizing combinatorial auctions. *Oper. Res.* 5, 63 (2015), 1000–1025.
- Shen, W., Tang, P., Zuo, S. Automated mechanism design via neural networks. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems*, 2019. Forthcoming.
- Thompson, D., Newman, N., Leyton-Brown, K. The positronic economist: A computational system for analyzing economic mechanisms. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, 2017, 720–727.
- Vickrey, W. Counterspeculation, auctions, and competitive sealed tenders. *J. Finance*, 16 (1961), 8–37.

Paul Dütting (duetting@google.com), Google Research, Zürich, Switzerland.

Zhe Feng (zhe_feng@g.harvard.edu), School of Engineering and Applied Sciences at Harvard University, Cambridge, MA, USA. This work is supported in part through a Google Ph.D. Fellowship.

Sai S. Ravindranath (saisr@g.harvard.edu), School of Engineering and Applied Sciences at Harvard University, Cambridge MA, USA.

David C. Parkes (parkes@eecs.harvard.edu), Harvard University, MA, USA. This work is supported in part through NSF award CCF-1841550.

Harikrishna Narasimhan (hnarasimhan@google.com), Google Research, Mountain View, CA, USA.

Copyright held by authors/owners. Publication rights licensed to ACM.

ACM Computing Surveys (CSUR)

2018 JOURNAL
IMPACT FACTOR:
6.131

Integration of computer science and engineering knowledge

ACM Computing Surveys (CSUR) publishes comprehensive, readable tutorials and survey papers that give guided tours through the literature and explain topics to those who seek to learn the basics of areas outside their specialties. These carefully planned and presented introductions are also an excellent way for professionals to develop perspectives on, and identify trends in, complex technologies.

For further information and to submit
your manuscript, visit csur.acm.org